

State space ARIMA for supply chain forecasting

Ivan Svetunkov^{a,*}, John E. Boylan^a

^a*Centre for Marketing Analytics and Forecasting
Lancaster University Management School, Lancaster, LA1 4YX, UK*

Abstract

ARIMA is seldom used in supply chains in practice. There are several reasons, not the least of which is the small sample size of available data, which restricts the usage of the model. Keeping in mind this restriction, we discuss in this paper a state space ARIMA model with a single source of error and show how it can be efficiently used in the supply chain context, especially in cases when only two seasonal cycles of data are available. We propose a new order selection algorithm for the model and compare its performance with the conventional ARIMA on real data. We show that the proposed model performs well in terms of both accuracy and computational time in comparison with other ARIMA implementations, which makes it efficient in the supply chain context.

Keywords: Forecasting, state space models, ARIMA, supply chain forecasting, order selection, model selection

1. Introduction

ARIMA has always been considered as a statistically sophisticated and complicated model. Although several forecasting competitions showed that simpler methods perform at least as well as statistically sophisticated methods and sometimes outperform ARIMA (Makridakis et al., 1982; Makridakis and Hibon, 1997, 2000; Athanasopoulos et al., 2011), its popularity among researchers has not declined over the years. ARIMA is considered to be a standard model in the statistical literature and is widely used for analytical

*Correspondance: Ivan Svetunkov, Department of Management Science, Lancaster University Management School, Lancaster, Lancashire, LA1 4YX, UK.

Email address: `i.svetunkov@lancaster.ac.uk` (Ivan Svetunkov)

9 derivations in the supply chain literature (an extensive review of supply chain
10 forecasting is given in Syntetos et al., 2016). Examples of the application of
11 ARIMA in a supply chain context include Kim et al. (2003), Wang et al.
12 (2010), Hosoda et al. (2008), Disney et al. (2006), Doganis et al. (2008),
13 Svetunkov and Petropoulos (2018), van Gils et al. (2017) and Dellino et al.
14 (2018).

15 Nevertheless ARIMA is not as widely used in practice as simpler meth-
16 ods, such as exponential smoothing and simple moving averages (Winklhofer
17 et al., 1996; Weller and Crone, 2012). The reason is the complexity of the
18 model. On the one hand it is not always simple to identify the appropriate
19 orders of ARIMA and estimate the model. On the other hand, it is much
20 harder to explain the model to supply chain managers than, for example, ex-
21 ponential smoothing. Furthermore, it is very common for companies working
22 in business to have small samples of data, because managers think that the
23 older data is not useful and not relevant to recent history. As a result com-
24 panies very often have at most 3 years of data. This makes seasonal ARIMA
25 models hard to build, because of estimation problems. Indeed, in order to
26 estimate the simplest conventional seasonal ARIMA, a forecaster needs at
27 least 3 years of data, where the first year is sacrificed for initialisation of the
28 model and the last two are needed for model fitting. Having less than 3 years
29 means that the model will overfit the second season and inevitably will pro-
30 duce poor forecasts. Furthermore, in order to include ARIMA in appropriate
31 forecasting evaluation against simpler forecasting methods, the sample needs
32 to be split into training and test sets. This further decreases the number of
33 observations available for estimation purposes, making conventional seasonal
34 ARIMA inapplicable.

35 Having limited data in the training set also means that parametric sta-
36 tistical tests may be inaccurate because of their low power on small samples.
37 This additional complication means that unit root tests and tests for season-
38 ality may be unreliable, which in turn leads to problems in the identification
39 of the correct order of ARIMA.

40 Finally, a typical forecasting task for the supply chain involves producing
41 forecasts for a large dataset with thousands of Stock Keeping Units (SKUs).
42 This means that the forecasting should be done automatically and fast, which
43 is not always the case for ARIMA models, because each time series has its own
44 structure, and the order of ARIMA needs to be selected individually. Order
45 selection for ARIMA is in general slow, because it either implies analysis of
46 Auto Correlation Functions (ACF and PACF), or applying several ARIMA

47 models of different orders to data and selecting the optimal one (using some
48 criterion).

49 All of this explains the lack of popularity of ARIMA models in applied
50 supply chain forecasting. At the same time interest in ARIMA models has
51 been recently rising, and overcoming the aforementioned limitations could
52 allow using the flexibility of ARIMA for supply chain forecasting. However,
53 this means that supply chain ARIMA should at least satisfy the following
54 requirements:

- 55 1. Order selection and model estimation should work with seasonal data
56 on small samples with at least two years of data;
- 57 2. Order selection should be done without statistical tests;
- 58 3. The order selection algorithm should be fast.

59 We propose using ARIMA in state space form with a Single Source of
60 Error (originally proposed in Snyder, 1985), which allows meeting all the
61 three requirements. First of all, state space models can be initialised in
62 period zero, which saves some observations and may increase the number of
63 degrees of freedom. Secondly, a state space model allows estimating ARIMA
64 using likelihood and applying model selection based on information criteria
65 for all the possible models without a need for hypotheses testing. The only
66 issue that needs to be addressed is the order selection algorithm, which should
67 be smart, choosing only those orders that are relevant to the data.

68 In this paper we discuss state space ARIMA and the methodology of order
69 selection and estimation of the model that satisfies all three requirements.
70 The proposed implementation of ARIMA can be efficiently applied to a wide
71 variety of data, and, as we show in the paper, performs well in terms of
72 forecasting accuracy, given the computational time restriction observed in
73 practical supply chains.

74 **2. State space ARIMA**

75 ARIMA in state space form has been known for at least 40 years. Harvey
76 and Phillips (1979) discuss a state space model with multiple sources of
77 errors (MSOE) underlying a general regression with ARMA errors. Pearlman
78 (1980) uses their finding and proposes a modification of the state space model
79 with a single source of error (SSOE). He points out that this model can be
80 used when the AR order is greater than or equal to MA order, but he does

81 not investigate the model further. Snyder (1985) analyses the SSOE state
 82 space model and its connection with ARIMA in more detail. He discusses
 83 several basic ARIMA models, showing how the model can be formulated
 84 using measurement and transition equations. Snyder et al. (2001) discuss
 85 ARIMA in state space form and demonstrate how the prediction intervals
 86 can be constructed for this model. Finally, a more detailed explanation of
 87 the connection between ARIMA and SSOE state space models is given in
 88 (Hyndman et al., 2008, pp. 173 - 174). We use their derivations as the basis
 89 for our model.

90 The general form of state space model with SSOE is (Hyndman et al.,
 91 2008):

$$\begin{aligned} y_t &= \mathbf{w}'\mathbf{v}_{t-1} + \epsilon_t \\ \mathbf{v}_t &= \mathbf{F}\mathbf{v}_{t-1} + \mathbf{g}\epsilon_t' \end{aligned} \quad (1)$$

92 where \mathbf{v}_t is the vector of states, ϵ_t is the error term (usually assumed to be
 93 distributed normally with zero mean and variance σ^2), \mathbf{F} is the transition
 94 matrix, \mathbf{w} is the measurement vector, \mathbf{w}' is the transposed \mathbf{w} and \mathbf{g} is the
 95 persistence vector. Hyndman et al. (2008) give general formulae, connecting
 96 ARIMA models with their state space counterparts. They derive the state
 97 space model for non-seasonal ARIMA without the constant term. Their
 98 derivations with minor modifications can be used in order to present the fol-
 99 lowing more general SARIMA(p, d, q)(P, D, Q) $_m$ model (where m is seasonal
 100 frequency) in state space form:

$$\phi_p(B)\delta_d(B)\Phi_P(B^m)\Delta_D(B^m)y_t = \theta_q(B)\Theta_Q(B^m)\epsilon_t + \beta, \quad (2)$$

101 where $\phi_p(B)$ is the non-seasonal AR, $\delta_d(B)$ is the non-seasonal difference,
 102 $\theta_q(B)$ is the non-seasonal MA, $\Phi_P(B^m)$ is the seasonal AR, $\Delta_D(B^m)$ is the
 103 seasonal differences and $\Theta_Q(B^m)$ is the seasonal MA polynomials, β is the
 104 constant term, which in the case of non-zero order of differences acts as drift
 105 and B is the backshift operator. We need to note that all the MA polynomials
 106 are used in our formulation with a plus sign, while the AR polynomials use
 107 the minus sign. So, for example, we formulate ARIMA(1,1,1) as:

$$(1 - \phi_1 B)(1 - B)y_t = (1 + \theta_1 B)\epsilon_t + \beta, \quad (3)$$

108 where ϕ_1 is AR(1) parameter and θ_1 is MA(1) parameter. By working models
 109 in this way we do not cause the confusion with signs of the coefficients.

110 In order to write ARIMA in state space form, the polynomials in the
 111 model (2) need to be expanded:

$$\left(1 - \sum_{j=1}^K \varphi_j B^j\right) y_t = \left(1 + \sum_{j=1}^K \eta_j B^j\right) \epsilon_t + \beta, \quad (4)$$

112 where φ_j and η_j are the values of the coefficients for AR and MA polynomials
 113 respectively and $K = \max(p+d+P+D, q+Q)$. The max term means that,
 114 for example, if $p+d+P+D > q+Q$, then all the η_j for $j > q+Q$ will be equal
 115 to zero. A similar property holds for the opposite situation. Regrouping the
 116 elements in (4) leads to:

$$y_t = \sum_{j=1}^K \varphi_j B^j y_{t-j} + \sum_{j=1}^K \eta_j B^j \epsilon_{t-j} + \beta + \epsilon_t. \quad (5)$$

117 After that the logic of derivation becomes exactly the same as in (Hyndman
 118 et al., 2008, pp. 173 - 174) with an exception for the first component of
 119 the state space model and an additional component for β . The state space
 120 ARIMA model proposed in this paper can be formulated in the following
 121 way:

$$\begin{aligned} y_t &= v_{1,t-1} + \epsilon_t \\ v_{j,t} &= \varphi_j v_{1,t-1} + v_{j+1,t-1} + v_{K+1,t-1} + (\varphi_j + \eta_j) \epsilon_t, & \text{for } j = 1 \\ v_{j,t} &= \varphi_j v_{1,t-1} + v_{j+1,t-1} + (\varphi_j + \eta_j) \epsilon_t, & \text{for } 1 < j \leq K \\ v_{K+1,t} &= v_{K+1,t-1}, \end{aligned} \quad (6)$$

122 where $v_{j,t}$ is the j -th component and $v_{K+1,0} = \beta$. Note that the first and the
 123 $K + 1$ components are calculated differently than in (Hyndman et al., 2008,
 124 pp. 173 - 174), because of the constant term β . If the constant is not needed
 125 for a time series, then $v_{K+1,0}$ can be set to zero, and the ARIMA model in
 126 state space form (6) becomes equivalent to the one in (Hyndman et al., 2008,
 127 p. 174). The model (6) can be written in the compact form (1), with:

$$\mathbf{v}_t = \begin{pmatrix} v_{1,t} \\ v_{2,t} \\ \vdots \\ v_{K,t} \\ v_{K+1,t} \end{pmatrix}, \mathbf{F} = \begin{pmatrix} \varphi_1 & 1 & 0 & \dots & 0 & 1 \\ \varphi_2 & 0 & 1 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ \varphi_K & 0 & 0 & \dots & 0 & 0 \\ 0 & 0 & 0 & \dots & 0 & 1 \end{pmatrix}, \mathbf{w} = \begin{pmatrix} 1 \\ 0 \\ \vdots \\ 0 \\ 0 \end{pmatrix}, \mathbf{g} = \begin{pmatrix} \varphi_1 + \eta_1 \\ \varphi_2 + \eta_2 \\ \vdots \\ \varphi_K + \eta_K \\ 0 \end{pmatrix}. \quad (7)$$

128 So ARIMA in state space form has $K + 1$ components if the constant term
 129 is not equal to zero. In cases with seasonal models the matrices in (7) can
 130 become large, especially if the seasonality lag m is large and the seasonal
 131 orders are high.

132 One of the advantages of the state space ARIMA model is that the ini-
 133 tialisation of the model (6) can be done on observation $t = 0$, which allows
 134 preserving observations for estimation purposes. The values of v_0 can be
 135 estimated in different ways, the most popular of which are optimisation and
 136 backcasting. We propose using the backcasting technique in order to preserve
 137 degrees of freedom and minimise the required computations, because then we
 138 do not need to estimate all the $K + 1$ initial values of the state vector; we only
 139 need to optimise the constant β which corresponds to the component $v_{K+1,0}$
 140 and the parameters of the ARMA. Still before constructing the model some
 141 preset values for the initial state vector are needed. In order to speed up the
 142 convergence to the true value of the initial vector v_0 , we use the following
 143 heuristics derived from the model (6) (see Appendix A):

$$\begin{aligned}
 v_{1,t-1} &= y_t, & \text{for } t = \{1, \dots, K\} \\
 v_{2,t-1} &= v_{1,t} - \varphi_1 y_t - v_{K+1,0}, & \text{for } t = \{1, \dots, K-1\} \\
 v_{j,t-1} &= v_{j-1,t} - \varphi_{j-1} y_t, & \text{for } 2 < j \leq K \text{ and } t = \{1, \dots, K-j+1\}
 \end{aligned} \tag{8}$$

144 In this way, we define $\frac{K(K+1)}{2}$ elements of the first K state vectors. After that
 145 the model (6) is applied to the data starting from the $t = 1$ until the last
 146 observation T in the sample. Then the reverse state space model is applied:

$$\begin{aligned}
 y_t &= \mathbf{w}'\mathbf{v}_{t+1} + \epsilon_t \\
 \mathbf{v}_t &= \mathbf{F}\mathbf{v}_{t+1} + \mathbf{g}\epsilon_t
 \end{aligned} \tag{9}$$

147 until the observation $t = 0$. Then a new initial value of the state vector is
 148 obtained and used in the construction of the model using (1). The procedure
 149 is repeated several times, refining the initial values of the state vector. In the
 150 implementation that we discuss in Section 4, three iterations are sufficient
 151 for the initial states to converge.

152 Having the state space model (6) also solves the problem with application
 153 of ARIMA to small samples. While in order to construct the conventional
 154 seasonal model it is necessary to have at least three seasonal cycles of data,
 155 the model (6) can be constructed even if only two seasonal cycles are avail-
 156 able. This is because the initialisation is done on the observation $t = 0$.
 157 For obvious reasons the estimates of the parameters on such a small sample

158 can be unreliable and the forecasts may be less accurate than they would
159 be on large samples, but at least some estimates and some forecasts can be
160 produced in this case.

161 The other important advantage of the model (6), is that all the possible
162 orders of the model can be compared directly with each other using infor-
163 mation criteria. Note that ARIMA models in the conventional form can be
164 compared with each other only for pre-specified differences, because taking
165 differences decreases the sample size, automatically leading to incomparable
166 values of information criteria. So there is no need to conduct preliminary
167 unit root tests in order to determine if the time series is stationary or not
168 with the state space ARIMA. There is also no need to test whether the se-
169 ries is seasonal or not, because this can be done by comparing seasonal and
170 non-seasonal ARIMA models in the state space form using an information
171 criterion.

172 However, taking into account that there are several possible orders in
173 seasonal ARIMA for each of the components of the model, the search of the
174 optimal order can become a cumbersome task. For example, if the maximum
175 orders of the model correspond to SARIMA(3,2,3)(2,1,2)_m for a fixed value of
176 m , then there are 864 potential models. Checking whether the constant β is
177 needed or not, doubles the number of models, giving a pool of 1728 SARIMA
178 models. In order to find a good model that would produce adequate forecasts
179 in a reasonable time, we need to use a smart algorithm for order selection.

180 **3. Order selection in state space ARIMA**

181 In order to select the most appropriate ARIMA, we propose using an in-
182 formation criterion. For example, the Akaike Information Criterion (Akaike,
183 1974) can be written as:

$$AIC = 2k - 2\ell, \quad (10)$$

184 where k is the number of estimated parameters and ℓ is the value of the
185 log-likelihood function extracted from the model.

186 We propose using the following stepwise order selection algorithm to allow
187 the selection of a good model for the data:

- 188 1. All the possible differences are checked with non-zero constant. This
189 includes seasonal and non-seasonal counterparts. In cases of non-zero
190 difference, the constant acts as a drift, allowing the capture of possible
191 trends in time series and model multiplicative seasonality.

192 2. The residuals of the best model on the step (1) are extracted. All
 193 possible types of seasonal and non-seasonal MA are checked. The order
 194 is selected via a modified information criterion, where the number of
 195 parameters is set to be equal to the sum of all the parameters estimated
 196 on the current and the previous steps:

$$\text{AIC}_2 = 2(k_1 + k_2) - 2\ell_2, \quad (11)$$

197 where the index in the subscript stands for the step in the algorithm,
 198 so that the number of parameters in AIC_2 is equal to sum of all the
 199 estimated parameters on step 2 and before. ℓ_2 is the value of the log-
 200 likelihood function for the model on step 2.

201 3. The residuals of the best model on the step (2) are extracted. All
 202 possible types of seasonal and non-seasonal AR orders are checked.
 203 The information criterion on this step uses the sum of all the estimated
 204 parameters on steps (1), (2) and (3), substituting $k_1 + k_2$ from (11) with
 205 $k_1 + k_2 + k_3$ and ℓ_2 with ℓ_3 , the value of log-likelihood function from
 206 the model on step 3.

207 4. The model of the selected orders is re-estimated on the original data in
 208 order to remove a potential bias in estimates of parameters.

209 5. The model (4) is compared with the same model without the constant.
 210 The model with the lowest information criterion is then selected for the
 211 forecasting purpose.

212 Note that if some other criterion is preferred, then the formula (11) should
 213 be substituted by the desired formula, preserving the number of estimated
 214 parameters and using the value of the log-likelihood function extracted for
 215 each specific step.

216 This algorithm allows for a substantial reduction in the pool of models.
 217 For example, in the case of $\text{SARIMA}(3,2,3)(2,1,2)_m$ only 31 models need to
 218 be checked instead of 1728. This does not guarantee that the selected model
 219 will have the lowest AIC among all the 1728 potential SARIMA models, but
 220 it gives a reasonable model, as will be demonstrated later in the paper, which
 221 should suffice for forecasting purposes.

222 In order to further decrease the pool of the potential models, higher orders
 223 of AR or MA can be checked before the lower orders. In this case when the
 224 higher order leads to the higher information criterion, then there is no need
 225 to check lower orders, meaning that they can be skipped altogether. For

226 example, if the true model is AR(1) and we first compare AR(0) and AR(3),
227 then the latter should have a lower information criterion, as the AR(3) model
228 includes the correct order as a first element $\varphi_1 y_{t-1}$. AR(2) in turn should
229 be better than AR(3) in terms of information criterion, because it does not
230 contain the redundant term $\varphi_3 y_{t-3}$, and finally AR(1) is expected to have
231 the lowest information criterion as it does not contain any redundant terms.
232 If for some data we find that AR(2) has greater information criterion than
233 AR(3), then the check of AR(1) can be skipped. This shortcut allows saving
234 computational time further by decreasing the pool of models.

235 4. Evaluation of state space ARIMA performance

236 In order to see how the state space ARIMA performs, we test it in a real
237 time series experiment.

238 The state space ARIMA with the described order selection algorithm
239 is implemented in `auto.ssarima()` function in `smooth` package version 2.1.1
240 for R (Svetunkov, 2017). This model is denoted as “SSARIMA” in the experi-
241 ment. The maximum order of the model was restricted to SARIMA(3,2,3)(2,1,2)_m.
242 This restriction is motivated by the following. The differences of the non-
243 seasonal part should not exceed 2 because this might cause over-differencing
244 with the corresponding loss of information (Box and Jenkins, 1976, p.175).
245 Similarly, there is no point in going beyond the first difference of the sea-
246 sonal part of the model. Given that we deal with short data, we restrict the
247 maximum seasonal orders of AR and MA to 2, which corresponds to two
248 years of data. Finally, the restriction of AR and MA to the maximum order
249 of 3 should be sufficient for such short data (this is similar to Hyndman and
250 Khandakar, 2008, who also investigated automatic model selection).

251 We have also applied state space ARIMA with optimised initials using
252 `auto.ssarima()` function (denoted “SSARIMA Opt”) in order to see the
253 influence of different initialisation techniques on forecasting accuracy.

254 In addition we used `auto.ssarima()` with backcasting and the switched
255 off mechanism of skipping orders (controlled by `workFast=FALSE` parameter),
256 discussed in the last paragraph of Section 3 in order to see if it improves the
257 performance of the model or not (this is denoted as “SSARIMA NSO”).

258 Furthermore, we have used the extensive search for state space ARIMA
259 (“SSARIMA Ext”), applying models with all the possible orders and selecting
260 the one with the lowest AIC. This took the most computational time, but
261 allowed us to evaluate the proposed algorithm of order selection.

262 Finally we have also used a benchmark in the experiment – conventional
 263 ARIMA implemented in `auto.arima()` function from `forecast` v8.4 package
 264 for R (Hyndman and Khandakar, 2008), denoted as “ARIMA”. This imple-
 265 mentation allows selecting between seasonal and non-seasonal models using
 266 information criteria, but the model itself is formulated in the conventional
 267 way.

268 In order to assess the accuracy of the proposed model, we use the data of
 269 an American retail company. This is typical supply chain data, containing
 270 4267 series with 36 monthly observations each. All the time series in the
 271 dataset can be categorised as shown in Table 1. The classification was done
 272 ex post, by applying the `auto.arima()` function to each of the time series
 273 and the whole 36 observations. We used the rule, according to which the
 274 time series is considered as seasonal, if seasonal AR, I or MA has non-zero
 275 order. If the resulting ARIMA model contained either a non-zero order of
 276 non-seasonal difference or a drift component, then the series was flagged as
 non-stationary. The categorisation in Table 1 is provided for information,

	Non-seasonal	Seasonal
Stationary	25.4%	23.5%
Non-stationary	17.4%	33.7%

Table 1: Categories of time series in the supply chain dataset.

277 showing the variety of different processes in the dataset, and it was not used
 278 for order selection or parameter evaluation.

280 In order to assess the accuracy of models, we withheld the last 9 observa-
 281 tions, which leaves 27 observations in the training set. This is a small sample
 282 from a conventional ARIMA perspective, but typical for supply chains and
 283 sufficient for simpler forecasting models. We do fixed origin evaluation, pro-
 284 ducing one to nine steps ahead forecasts, and then calculating the following
 285 error measures:

1. MPE – Mean Percentage Error, which assesses bias of forecasts:

$$\text{MPE} = \frac{1}{h} \sum_{j=1}^h \frac{e_{t+j}}{y_{t+j}},$$

286 where e_{t+j} is the j -steps ahead forecast error, and $h = 9$ is the fore-
 287 casting horizon, for this and all the other error measures.

2. MAPE – Mean Absolute Percentage Error, which assesses accuracy of forecasts and is commonly used in practice:

$$\text{MAPE} = \frac{1}{h} \sum_{j=1}^h \frac{|e_{t+j}|}{y_{t+j}}.$$

288 This is considered by many forecasters as a biased error measure as it
289 encourages under-forecasting (Makridakis, 1993).

3. MASE – Mean Absolute Scaled Error, measure proposed by Hyndman and Koehler (2006), which is less biased than MAPE:

$$\text{MASE} = \frac{\frac{1}{h} \sum_{j=1}^h |e_{t+j}|}{\frac{1}{t-1} \sum_{i=2}^t |y_i - y_{i-1}|};$$

4. sMAE – scaled Mean Absolute Error by Petropoulos and Kourentzes (2015), which is similar to MASE, but has easier interpretation, close to the one of MAPE:

$$\text{sMAE} = \frac{\frac{1}{h} \sum_{j=1}^h |e_{t+j}|}{\bar{y}},$$

290 where $\bar{y} = \frac{1}{t} \sum_{i=1}^t y_i$.

5. ARMAE – Average Relative Mean Absolute Error from Davydenko and Fildes (2013) which was shown to be the least biased error measure among (2) – (5):

$$\text{ARMAE} = \frac{\frac{1}{h} \sum_{j=1}^h |e_{1,t+j}|}{\frac{1}{h} \sum_{j=1}^h |e_{2,t+j}|},$$

291 where $e_{1,t+j}$ is the j -steps ahead forecast error of the model under
292 consideration and $e_{2,t+j}$ is the j -steps ahead forecast error of the Naïve
293 method. Note that when ARMAE is aggregated over all the series, the
294 geometric mean is used instead of arithmetic.

295 The error measures have been calculated for each separate time series.
296 After that the mean and the median values of each error measure across all
297 the series have been calculated. The results are presented in Tables 2 and 3.
298 The best values in the tables are shown in boldface; the second best values
299 are in italic.

Model	MPE	MAPE	MASE	sMAE	ARMAE
ARIMA	-18.2	49.4	<i>119.6</i>	41.5	91.0
SSARIMA	-15.4	<i>48.4</i>	119.2	<i>41.3</i>	90.0
SSARIMA NSO	<i>-14.8</i>	48.1	119.2	<i>41.3</i>	<i>89.9</i>
SSARIMA Opt	-15.1	48.8	120.2	41.2	89.9
SSARIMA Ext	-11.4	50.9	126.5	44.0	95.2

Table 2: Mean error measures (percentages).

300 As can be seen from Table 2, SSARIMA with backcasting and the pro-
301 posed order selection method performs better or at least not worse than more
302 complicated SSARIMA models, including the one with the extensive search:
303 the differences in performance of SSARIMA with the other versions of the
304 model are very small.

305 Similar conclusions can be drawn from the analysis of Table 3, where
306 SSARIMA performed slightly better than the other models in terms of MASE
307 and ARMAE. Note that SSARIMA performed better than the conventional
308 ARIMA across all measures. This can be explained by the ability of the
309 former to better identify seasonality on small samples.

Model	MPE	MAPE	MASE	sMAE	ARMAE
ARIMA	-4.1	33.5	103.9	34.8	97.4
SSARIMA	<i>-1.9</i>	<i>32.5</i>	100.0	<i>34.4</i>	92.9
SSARIMA NSO	-2.1	32.7	<i>100.6</i>	34.4	<i>93.2</i>
SSARIMA Opt	-1.0	32.4	102.8	34.0	96.7
SSARIMA Ext	-0.5	35.0	105.0	35.8	95.8

Table 3: Median error measures (percentages).

310 It is worth noting that the extensive search of the optimal order does
311 not improve upon the accuracy of SSARIMA model – although the order
312 selected by SSARIMA is not optimal in the sense of AIC, it performs better
313 in terms of forecasting accuracy. Furthermore, the optimisation procedure
314 does not bring significant improvements and the SSARIMA NSO performs
315 slightly worse than SSARIMA in many cases. In addition, the SSARIMA
316 with the optimised initials outperforms SSARIMA with backcasting in some
317 cases, but it does not demonstrate substantial improvement.

318 In order to further investigate the performance of SSARIMA versus ARIMA,
 319 we summarise the ARMAE error measures for the four categories from Table
 320 1. The results for the other error measures look similar, so we have decided
 321 to focus on ARMAE, as it is the least biased error measures of the ones in
 322 our pool (Davydenko and Fildes, 2013). These values are presented in Table
 323 4.

Series type	ARIMA	SSARIMA
Non Seasonal, Stationary	0.798	0.824
Non Seasonal, Non Stationary	1.050	1.155
Seasonal, Stationary	0.872	0.817
Seasonal, Non Stationary	0.961	0.906

Table 4: ARMAE of ARIMA and SSARIMA for different categories of the data.

324 It can be noted from Table 4, that while ARIMA performs better than
 325 SSARIMA on non-seasonal time series, SSARIMA is much better on the
 326 seasonal data, thus showing the improvement in the overall forecasting ac-
 327 curacy. In fact, it seems that SSARIMA overfits the non-seasonal data,
 328 selecting wrongly the seasonal orders, while ARIMA underfits the seasonal
 329 data, not selecting the necessary orders. Given the prevalence of seasonal
 330 data in the dataset (57.2% according to Table 1), the summary value of the
 331 ARMAE for SSARIMA is lower than that of ARIMA.

332 Finally, Table 5 summarises the computational time for each of the models
 333 for the whole dataset (calculated in serial on Intel Core i7 of 5th generation).

Model	Time in minutes
ARIMA	39.82
SSARIMA	47.86
SSARIMA NSO	309.67
SSARIMA Opt	289.68
SSARIMA Ext	17,989.72

Table 5: Time of computation for each model in minutes for all 4267 series.

334 Although SSARIMA could not outperform ARIMA in terms of time, the
 335 difference in their performance is not large. At the same time SSARIMA with

336 the proposed algorithm produced forecasts faster than any other implemen-
337 tation, and in a practical time. As expected, SSARIMA with backcasting
338 and the new order selection performed much faster than other SSARIMA
339 algorithms. Note that the extensive search took almost 18 thousand min-
340 utes of computational time, which is equivalent to 300 hours or 12.5 days.
341 Nevertheless, it performed worse than the faster algorithms. So, although
342 SSARIMA does not necessarily beat other models in forecasting accuracy, it
343 is much more efficient and faster than its competitors. Taking into account
344 the accuracy of the state space ARIMA and its speed of calculation, it can
345 be concluded that the model in the proposed form can efficiently be used in
346 a supply chain context, especially for seasonal data.

347 5. Conclusions

348 ARIMA is seldom used in a supply chain context because of the limita-
349 tions of the data and general complexity of the model. We have discussed
350 the state space form of ARIMA with a single source of error and showed that
351 it overcomes some of the limitations of the conventional ARIMA. We have
352 shown that the state space ARIMA simplifies some of the steps in forecasting
353 and can be used even on data with a short history.

354 All of the above allows using seasonal ARIMA on small samples, contain-
355 ing at least 2 seasonal cycles, something that ARIMA in the conventional
356 form cannot do. In addition the state space form permits comparing differ-
357 ent models directly using information criteria, because they can be initialised
358 in the zero period, making sample sizes for models with different orders of
359 differences the same.

360 We have also proposed an algorithm of order selection for state space
361 ARIMA, which substantially decreases the pool of models under consider-
362 ation. This algorithm does not employ hypothesis testing, an important
363 feature in cases of small samples, which are very common in a supply chain
364 context. We tested the state space ARIMA with the proposed order selec-
365 tion algorithm on supply chain data and showed that it outperforms the
366 implementation of the conventional ARIMA from `forecast` package for R
367 in terms of accuracy and that it works fast. It seems to perform especially
368 well on seasonal data. Further research is ongoing to find improvements in
369 the algorithm for non-seasonal data.

370 Furthermore, the practicality of our proposed approach is evidenced by its
371 introduction into commercial software by the *Demand Works* company. The

372 new software module, called ARIMA, is based on the SSARIMA model dis-
 373 cussed in this paper. However, it has been subject to several modifications
 374 and adjustments which cannot be disclosed because of confidentiality rea-
 375 sons. Nevertheless, we can report that the implemented SSARIMA module
 376 demonstrates further improvements in the accuracy and significant reduction
 377 in computational times in comparison with the implementations discussed in
 378 this paper. Furthermore, *Demand Works* software is used by over 400 corpo-
 379 rations, demonstrating that the approach discussed in this paper has reason-
 380 able commercial applicability. In summary, we can conclude that ARIMA in
 381 state space form is a practical and efficient option for supply-chain forecast-
 382 ing and, indeed, for any context, where the historical time series is limited
 383 with few complete seasonal cycles.

384 The focus of this paper was on state-space ARIMA applied to supply
 385 chain data. However, this is not the only possible area of application, and
 386 we think that developing and exploring the efficient algorithms for ARIMA
 387 application in other business contexts is an interesting direction for future
 388 research. This means that as a future work, the state-space ARIMA should
 389 be tested on other datasets and compared with other popular forecasting
 390 methods. Finally, another interesting direction for future work would be to
 391 compare the performance of the proposed approach with the other approaches
 392 in terms of inventory measures, such as service level and costs of stocking,
 393 similar to the analysis by Syntetos and Boylan (2006).

394 **Acknowledgement**

395 We would like to thank *Demand Works* and especially Bill Tonetti and
 396 Eric Townson for their support and cooperation in testing of the models and
 397 their help in implementation of the ARIMA module in Smoothie.

398 **Appendix A. Derivation of initial values of state vector**

399 We assume that $\epsilon_j = 0$ for each $j = 1, \dots, K$, which gives us K estimates
 400 of the first component of the vector based on the measurement equation in
 401 (6):

$$\begin{aligned}
 v_{1,0} &= y_1 \\
 v_{1,1} &= y_2 \\
 &\vdots \\
 v_{1,K-1} &= y_K
 \end{aligned}
 \tag{A.1}$$

402 This means that the state space model (6) simplifies to:

$$\begin{aligned}
 y_t &= v_{1,t-1} \\
 v_{1,t} &= \varphi_1 v_{1,t-1} + v_{2,t-1} + v_{K+1,t-1}, & \text{for } j = 1 \\
 v_{j,t} &= \varphi_j v_{1,t-1} + v_{j+1,t-1}, & \text{for } 1 < j \leq K \\
 v_{K+1,t} &= v_{K+1,t-1}
 \end{aligned} \tag{A.2}$$

403 Every $j + 1$ component for $1 < j \leq K$ in (A.2) can be expressed the following
 404 way:

$$v_{j+1,t-1} = v_{j,t} - \varphi_j v_{1,t-1}, \tag{A.3}$$

405 meaning that it can be expressed using the values of the previous component
 406 and the very first one. The second component is expressed as:

$$v_{2,t-1} = v_{1,t} - \varphi_1 v_{1,t-1} - v_{K+1,t-1}. \tag{A.4}$$

407 Substituting values from (A.1) into (A.3) and (A.4) leads to the following
 408 system:

$$\begin{aligned}
 v_{1,t-1} &= y_t, & \text{for } t = \{1, \dots, K\} \\
 v_{2,t-1} &= v_{1,t} - \varphi_1 y_t - v_{K+1,t-1}, & \text{for } t = \{1, \dots, K - 1\} \\
 v_{j,t-1} &= v_{j-1,t} - \varphi_{j-1} y_t, & \text{for } 2 < j \leq K \text{ and } t = \{1, \dots, K - j + 1\}
 \end{aligned} \tag{A.5}$$

409 So the procedure of the initialisation of the state vector of state space ARIMA
 410 is iterative, the components are defined one after another, starting from the
 411 first and finishing with the K -th. The value of $K + 1$ component in this case
 412 is defined by the optimiser.

413 References

- 414 Akaike, H., 1974. A new look at the statistical model identification. IEEE
 415 Transactions on Automatic Control 19 (6), 716–723.
- 416 Athanasopoulos, G., Hyndman, R. J., Song, H., Wu, D. C., 2011. The tourism
 417 forecasting competition. International Journal of Forecasting 27 (3), 822–
 418 844.
- 419 Box, G., Jenkins, G., 1976. Time series analysis: forecasting and control.
 420 Holden-day, Oakland, California.

- 421 Davydenko, A., Fildes, R., 2013. Measuring Forecasting Accuracy: The Case
422 Of Judgmental Adjustments To SKU-Level Demand Forecasts. *International*
423 *Journal of Forecasting* 29 (3), 510–522.
- 424 Dellino, G., Laudadio, T., Mari, R., Mastronardi, N., Meloni, C., 2018. A
425 reliable decision support system for fresh food supply chain management.
426 *International Journal of Production Research* 56 (4), 1458–1485.
- 427 Disney, S., Farasyn, I., Lambrecht, M., Towill, D., de Velde, W. V., 2006.
428 Taming the bullwhip effect whilst watching customer service in a single
429 supply chain echelon. *European Journal of Operational Research* 173 (1),
430 151 – 172.
- 431 Doganis, P., Aggelogiannaki, E., Sarimveis, H., 2008. A combined model
432 predictive control and time series forecasting framework for production-
433 inventory systems. *International Journal of Production Research* 46 (24),
434 6841–6853.
- 435 Harvey, A. C., Phillips, G. D., 1979. Maximum Likelihood Estimation
436 of Regression Models with Autoregressive-Moving Average Disturbances.
437 *Biometrika* 66 (1), 49–58.
- 438 Hosoda, T., Naim, M. M., Disney, S. M., Potter, A., 2008. Is there a benefit to
439 sharing market sales information? linking theory and practice. *Computers*
440 *& Industrial Engineering* 54 (2), 315 – 326.
- 441 Hyndman, R. J., Khandakar, Y., 2008. Automatic time series forecasting:
442 the forecast package for R. *Journal of Statistical Software* 27 (3), 1–22.
- 443 Hyndman, R. J., Koehler, A. B., 2006. Another look at measures of forecast
444 accuracy. *International Journal of Forecasting* 22 (4), 679–688.
- 445 Hyndman, R. J., Koehler, A. B., Ord, J. K., Snyder, R. D., 2008. *Forecasting*
446 *with Exponential Smoothing*. Springer Berlin Heidelberg.
- 447 Kim, J. S., Shin, K. Y., Ahn, S. E., 2003. A multiple replenishment contract
448 with ARIMA demand processes. *Journal of the Operational Research So-*
449 *ciety* 54 (11), 1189–1197.
- 450 Makridakis, S., 1993. Accuracy concerns measures : theoretical and practical
451 concerns. *International Journal of Forecasting* 9, 527–529.

- 452 Makridakis, S., Andersen, A. P., Carbone, R., Fildes, R., Hibon, M.,
453 Lewandowski, R., Newton, J., Parzen, E., Winkler, R. L., 1982. The ac-
454 curacy of extrapolation (time series) methods: Results of a forecasting
455 competition. *Journal of Forecasting* 1 (2), 111–153.
- 456 Makridakis, S., Hibon, M., 1997. ARMA models and the Box-Jenkins
457 methodology. *Journal of Forecasting* 16, 147–163.
- 458 Makridakis, S., Hibon, M., 2000. The M3-Competition: results, conclusions
459 and implications. *International Journal of Forecasting* 16, 451–476.
- 460 Pearlman, J. G., 1980. An algorithm for the exact likelihood of a mixed
461 autoregressive moving average process. *Biometrika* 67 (1), 232–233.
- 462 Petropoulos, F., Kourentzes, N., 2015. Forecast combinations for intermittent
463 demand. *Journal of the Operational Research Society* 66 (6), 914–924.
- 464 Snyder, R. D., 1985. Recursive Estimation of Dynamic Linear Models. *Jour-
465 nal of the Royal Statistical Society, Series B (Methodological)* 47 (2), 272–
466 276.
- 467 Snyder, R. D., Ord, J. K., Koehler, A. B., 2001. Prediction Intervals for
468 ARIMA Models. *Journal of Business & Economic Statistics* 19 (2), 217–
469 225.
- 470 Svetunkov, I., 2017. smooth: Forecasting Using Smoothing Functions. R
471 package version 2.1.1.
472 URL <https://github.com/config-i1/smooth>
- 473 Svetunkov, I., Petropoulos, F., 2018. Old dog, new tricks: a modelling view
474 of simple moving averages. *International Journal of Production Research*
475 56 (18), 6034–6047.
- 476 Syntetos, A. A., Babai, Z., Boylan, J. E., Kolassa, S., Nikolopoulos, K., 2016.
477 Supply Chain Forecasting: Theory, Practice, their Gap and the Future.
478 *European Journal of Operational Research* 252, 1–26.
- 479 Syntetos, A. A., Boylan, J. E., 2006. On the stock control performance of
480 intermittent demand estimators. *International Journal of Production Eco-
481 nomics* 103 (1), 36–47.

- 482 van Gils, T., Ramaekers, K., Caris, A., Cools, M., 2017. The use of time
483 series forecasting in zone order picking systems to predict order pickers'
484 workload. *International Journal of Production Research* 55 (21), 6380–
485 6393.
- 486 Wang, S. J., Huang, C. T., Wang, W. L., Chen, Y. H., 2010. Incorporating
487 ARIMA forecasting and service-level based replenishment in RFID-enabled
488 supply chain. *International Journal of Production Research* 48 (9), 2655–
489 2677.
- 490 Weller, M., Crone, S., 2012. *Supply Chain Forecasting: Best Practices and*
491 *Benchmarking Study*. Tech. rep., Lancaster Centre for Forecasting.
- 492 Winklhofer, H., Diamantopoulos, A., Witt, S. F., 1996. Forecasting practice:
493 A review of the empirical literature and an agenda for future research.
494 *International Journal of Forecasting* 12 (2), 193–221.